

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
10 May 2001 (10.05.2001)

PCT

(10) International Publication Number
WO 01/33863 A1

(51) International Patent Classification⁷: H04N 7/36,
5/14, G06F 17/30

Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
ELENBAAS, Jan, H.; Prof. Holstlaan 6, NL-5656 AA
Eindhoven (NL).

(21) International Application Number: PCT/EP00/10161

(22) International Filing Date: 13 October 2000 (13.10.2000)

(74) Agent: WHITE, Andrew, G.; Internationaal Octrooibu-
reau B.V., Prof. Holstlaan 6, NL-5656 AA Eindhoven
(NL).

(25) Filing Language: English

(81) Designated States (*national*): JP, KR.

(26) Publication Language: English

(30) Priority Data:
09/433,258 4 November 1999 (04.11.1999) US

(84) Designated States (*regional*): European patent (AT, BE,
CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC,
NL, PT, SE).


(71) Applicant: KONINKLIJKE PHILIPS ELECTRON-
ICS N.V. [NL/NL]; Groenewoudseweg 1, NL-5621 BA
Eindhoven (NL).

Published:
— With international search report.

(72) Inventors: MCGEE, Thomas; Prof. Holstlaan 6,
NL-5656 AA Eindhoven (NL). DIMITROVA, Nevenka;

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: SIGNIFICANT SCENE DETECTION AND FRAME FILTERING FOR A VISUAL INDEXING SYSTEM USING DYNAMIC THRESHOLD



(57) Abstract: A video indexing system analyzes contents of source video and develops a visual table of contents using selected images. A system for detecting significant scenes detects video cuts from one scene to another, and static scenes based on DCT coefficients and macroblocks. A keyframe filtering process filters out less desired frames including, for example, unicolor frames, or those frames having a same object as a primary focus or one primary focuses. Commercials may also be detected and frames of commercials eliminated. The significant scenes and static scenes are detected based on a threshold which is set based on the category of the video.

WO 01/33863 A1

Significant scene detection and frame filtering for a visual indexing system using dynamic threshold

The present invention is related to an apparatus which detects significant scenes of a source video and selects keyframes to represent each detected significant scene. The present invention additionally filters the selected keyframes and creates a visual index or a visual table of contents based on remaining keyframes.

- 5 Users will often record home videos or record television programs, movies, concerts, sports events, etc. on a tape for later or repeated viewing. Often, a video will have varied content or be of great length. However, a user may not write down what is on a recorded tape and may not remember what she recorded on a tape, DVD or other medium or where on a tape particular scenes, movies, or events are recorded. Thus, a user may have to sit and view an
10 entire tape to remember what is on the tape.

Video content analysis uses automatic and semi-automatic methods to extract information that describes contents of the recorded material. Video content indexing and analysis extracts structure and meaning from visual cues in the video. Generally, a video clip is taken from a TV program or a home video.

- 15 In U.S. Serial No. 08/867140, of which the present application is a continuation in part thereof, a method and device is described which detects scene changes or "cuts" in the video. At least one frame between detected cuts is then selected as a key frame to create a video index. In order to detect scene changes a first frame is selected and then a subsequent frame is compared to the first frame and a difference calculation is made which
20 represents the content difference between the two frames. The result of this difference calculation is then compared to a universal threshold or thresholds which is/are used for all categories of video. If the difference is above the universal threshold(s) it is determined that a scene change has occurred.

- 25 In US 08/867140 a universal threshold(s) is/are chosen which is/are optimal for all types of video. The problem with such an application is that a visual index of a video which contains high action, such as an action movie, will be quite large, whereas a visual index of a video with little action, such as the news will be quite small. This is because in a high action movie, where objects are moving across a scene, the content difference between two consecutive frames may be large. In such a case, comparing the content difference to a

universal threshold will result in a "cut" being detected even though the two frames may be within the same scene. If there are more perceived cuts or scene changes then there will be more key frames and vice versa. Accordingly an action movie ends up having far too many key frames to represent the movie.

5

Accordingly it is an object of the invention to provide a system which will create a visual index for a video source which was previously recorded or while being recorded, which is useable and more accurate in selecting significant keyframes by varying the number of keyframes chosen from a video based on the category of the video.

10

The present invention further presents a video analysis system supporting visual content extraction for source video which may include informative and/or entertainment programs such as news, serials, weather, sports or any type of home recorded video, broadcast material, prerecorded material etc.

15

The present invention further presents a new apparatus for video cut detection, static scene detection, and keyframe filtering to provide for more useable visual images in the visual index by comparing two frames of video and determining whether the differences between the frames are above or below a certain threshold, the threshold being dependent on the category of the video. If the differences are above the selected threshold then this point in the video is determined to be a scene "cut". The frames within the scene cuts can then be filtered to select key frames for a video index. If the differences are below a static scene threshold for a number of frames it is determined to be a static sequence.

20

The present invention may detect significant scenes and select keyframes of source video, based on calculations using DCT coefficients and comparisons to various thresholds where the thresholds vary based on the category of video, e.g. action, news, music, or even the category of portions of a video.

25

Additionally it is an object of the invention to calculate these thresholds based on a video category provided by electronic program guides, or alternatively instead of calculating the thresholds the electronic program guide provides the thresholds themselves.

30

Furthermore it is an object of the invention to enter these thresholds manually. It is even another object of the invention to provide the thresholds in the encoded video.

For a better understanding of the invention, its operating advantages and specific objects attained by its use, reference should be had to the accompanying drawings

and descriptive matter in which there are illustrated and described the preferred embodiments of the invention.

5 Figure 1 illustrates a video archival process;
 Figures 2A and 2B are block diagrams of devices used in creating a visual
index;
 Figure 3 illustrates a frame, a macroblock, and several blocks;
 Figure 4 illustrates several DCT coefficients of a block;
10 Figure 5 illustrates a macroblock and several blocks with DCT coefficients;
 Figures 6A and 6B illustrate a procedure for keyframe filtering;
 Figure 7 is illustrates a macroblock and its relationship to a block signature;
 Figure 8 illustrates a video retrieval process;
 Figure 8a illustrates a video keyframe system;
15 Figure 9 illustrates a threshold detecting device; and
 Figure 10 illustrates a threshold setting device.

The basic operation of a visual indexing system is first explained:

20 A visual index may be created on a pre-existing tape (or file, DVD, disks, etc.)
or while recording on a new tape. For both tapes a predetermined portion at a selected area on
the tape (if available), in this example, the beginning for ease of use, to allow a visual index
to be created or the index could be stored on an entirely different medium such as a hard disk
etc. For the present example, thirty seconds of "blank" or overwritable tape is desired. For a
25 file, the selected area for the visual index may occur anywhere in the file, and may be
reserved by a system automatically or manually selected by a user.

 The visual index may include visual images, audio, text or any combination
thereof. For the present example, visual images and text are provided. To create and use the
visual index, a video content indexing process is performed.

30 Two phases exist in the video content indexing process: archival and retrieval.
During the archival process, video content is analyzed during a video analysis process and a
visual index is created. In the video analysis process, automatic significant scene detection
and keyframe selection occur. Significant scene detection is a process of identifying scene
changes, i.e., "cuts" (video cut detection or segmentation detection) and identifying static

scenes (static scene detection). For each scene, a particular representative frame called a keyframe is extracted. A keyframe filtering and selection process is applied to each keyframe of source video, such as a video tape, to create a visual index from selectively chosen key frames. Reference is to a source tape although clearly, the source video may be from a file, disk, DVD, other storage means or directly from a transmission source (e.g., while recording a home video).

In video tape indexing, an index is generally stored on the source tape, CD etc. In video indexing an MPEG 1, MPEG 2, MPEG 4, Motion JPEG file or any other video file from a Video CD, DVD, or other storage device, or from a broadcast stream, the index may be stored on a hard disk, or other storage medium.

A video archival process is shown in Figure 1 for a source tape with previously recorded source video, which may include audio and/or text, although a similar process may be followed for other storage devices with previously saved visual information, such as an MPEG file. In this process, a visual index is created based on the source video. A second process, for a source tape on which a user intends to record, creates a visual index simultaneously with the recording.

Figure 1 illustrates an example of the first process (for previously recorded source tape) for a video tape. In step 101, the source video is rewound, if required, by a playback/recording device such as a VCR. In step 102, the source video is played back. Signals from the source video are received by a television, a VCR or other processing device. In step 103, a media processor in the processing device or an external processor, receives the video signals and formats the video signals into frames representing pixel data (frame grabbing).

In step 104, a host processor separates each frame into blocks, and transforms the blocks and their associated data to create DCT (discrete cosine transform) coefficients; performs significant scene detection and keyframe selection; and builds and stores keyframes as a data structure in a memory, disk or other storage medium. In step 105, the source tape is rewound to its beginning and in step 106, the source tape is set to record information. In step 107, the data structure is transferred from the memory to the source tape, creating the visual index. The tape may then be rewound to view the visual index.

The above process is slightly altered when a user wishes to create a visual index on a tape while recording. Instead of steps 101 and 102, as shown in step 112 of Figure 1, the frame grabbing process of step 103 occurs as the video (film, etc.) is being recorded.

Additionally, if the tape, or file, is not completely recorded on at one time, a partially created video index could be saved on the tape, file, etc. or could be saved in a tape memory for later additions.

Steps 103 and 104 are more specifically illustrated in Figures 2A and 2B.

- 5 Video exists either in analog (continuous data) or digital (discrete data) form. The present example operates in the digital domain and thus uses digital form for processing. The source video or video signal is thus a series of individual images or video frames displayed at a rate high enough (in this example 30 frames per second) so the displayed sequence of images appears as a continuous picture stream. These video frames may be uncompressed (NTSC or
10 raw video) or compressed data in a format such as MPEG, MPEG 2, MPEG 4, Motion JPEG or such.

- The information in an uncompressed video is first segmented into frames in a media processor 202, using a frame grabbing technique such as present on the Intel® Smart Video Recorder III. Although other frame sizes are available, in this example shown in
15 Figure 3, a frame 302 represents one television, video, or other visual image and includes 352 x 240 pixels.

- The frames 302 are each broken into blocks 304 of, in this example, 8 x 8 pixels in the host processor 210 (Figure 2A). Using these blocks 304 and a popular broadcast standard, CCIR-601, a macroblock creator 206 (Figure 2A) creates luminance blocks and
20 averages color information to create chrominance blocks. The luminance and chrominance blocks form a macroblock 308. In this example, 4:2:0 is being used although other formats such as 4:1:1 and 4:2:2 could easily be used by one skilled in the art. In 4:2:0, a macroblock 308 has six blocks, four luminance, Y1, Y2, Y3, and Y4; and two chrominance Cr and Cb, each block within a macroblock being 8x8 pixels.

- 25 The video signal may also represent a compressed image using a compression standard such as Motion JPEG (Joint Photographic Experts Group) and MPEG (Motion Pictures Experts Group). If the signal is instead an MPEG or other compressed signal, as shown in Figure 2B the MPEG signal is broken into frames using a frame or bitstream parsing technique by a frame parser 205. The frames are then sent to an entropy decoder 214
30 in the media processor 203 and to a table specifier 216. The entropy decoder 214 decodes the MPEG signal using data from the table specifier 216, using, for example, Huffman decoding, or another decoding technique.

The decoded signal is next supplied to a dequantizer 218 which dequantizes the decoded signal using data from the table specifier 216. Although shown as occurring in

the media processor 203, these steps (steps 214-218) may occur in either the media processor 203, host processor 211 or even another external device depending upon the devices used.

Alternatively, if a system has encoding capability (in the media processor, for example) that allows access at different stages of the processing, the DCT coefficients could be delivered directly to the host processor. In all these approaches, processing may be performed in up to real time.

In step 104 of Figure 1, the host processor 210, which may be, for example, an Intel® Pentium™ chip or other multiprocessor, a Philips® Trimedia™ chip or any other multimedia processor; a computer; an enhanced VCR, record/playback device, or television; or any other processor, performs significant scene detection, key frame selection, and building and storing a data structure in an index memory, such as, for example, a hard disk, file, tape, DVD, or other storage medium.

Significant Scene Detection: For automatic significant scene detection, the present invention attempts to detect when a scene of a video has changed or a static scene has occurred. A scene may represent one or more related images. In significant scene detection, two consecutive frames are compared and, if the frames are determined to be significantly different, a scene change is determined to have occurred between the two frames; and if determined to be significantly alike, processing is performed to determine if a static scene has occurred. In a static type scene, such as a news broadcast, a lesser change between two frames may indicate a scene change. In the same action move the same change between two frames may not be a scene change but instead it may only be an object moving across the scene.

From each scene, one or more keyframes is extracted to represent the scene. Typically, current theory proposes using the first video (visual) frame in a scene. However, in many cases, the main subject or event in a scene appears after a camera zoom or pan. Additionally, current theories typically do not detect scenes that remain constant for some length of time (Static Scenes). However, based on the length of time spent on that scene, from a videographer's, director's, etc. point of view this may have been an important scene.

Each of the present methods uses comparisons of DCT (Discrete Cosine Transform) coefficients. There are other methods that can be used such as histogram differences, wavelets, comparisons based on pixels, edges, fractals, other DCT methods or any other method which uses thresholds to compare frames. First, each received frame 302 is

processed individually in the host processor 210 to create macroblocks 308. The host processor 210 processes each macroblock 308 which contains spatial information, using a discrete cosine transformer 220 to extract DCT coefficients and create the six 8 x 8 blocks 440 (Figure 4) of DCT coefficients.

When the video signal received is in compressed video format such as MPEG, the DCT coefficients may be extracted after dequantization and need not be processed by a discrete cosine transformer. Additionally, as previously discussed, DCT coefficients may be automatically extracted depending upon the devices used.

The DCT transformer provides each of the blocks 440 (Figure 4), Y1, Y2, Y3, Y4, Cr and Cb with DCT coefficient values. According to this standard, the uppermost left hand corner of each block contains DC information (DC value) and the remaining DCT coefficients contain AC information (AC values). The AC values increase in frequency in a zig-zag order from the right of the DC value, to the DCT coefficient just beneath the DC value, as partially shown in Figure 4.

The present invention may use several different significant scene detection methods, all of which use the DCT coefficients for the respective block. The host processor 210 further processes each of the macroblocks using at least one of the following methods in the significant scene processor 230 (Figure 2A).

In the methods to follow, processing is limited to the DC values to more quickly produce results and limit processing without a significant loss in efficiency; however, clearly one skilled in the art could process all of the DCT coefficients for each of the macroblocks. By looping through each block using these steps, all the DCT coefficients could be analyzed, although this would affect time needed for processing.

Method One:

$$\text{SUM}[i] = \sum_{k,j} \text{ABS}(\text{DCT1}_{k,j}[i] - \text{DCT2}_{k,j}[i])$$

where:

k is the number of macroblocks in width of a frame, k = 1 to Frame-width/16,

j is the number of macroblocks in height of a frame, j = 1 to Frame-height/16,

i is the number of blocks in a macroblock, i = 1 to number of blocks in macroblock,

DCT1_{k,j} and DCT2_{k,j} are DCT coefficients for the specified macroblock for a previous and a current video frame, respectively, as illustrated in Figure 5 and

ABS is an absolute value function.

In this example, for a 352 x 240 pixel frame, $k = 1$ to 22, $j = 1$ to 15, and $i = 1$ to 6. In this method and the following methods, the macroblock width of a frame or the macroblock height of a frame will be an even integer, since if the frame sizes are not evenly divisible, the frame size is scaled to fit during processing.

5 Method one differentiates between each of the blocks (four luminance and two chrominance blocks) of the macroblocks. In this method, DC values for each luminance and chrominance block in a current macroblock from the current video frame are respectively subtracted from a corresponding DC value for a corresponding block in the previous video frame. Separate sums of differences, SUM [I], are kept for each luminance and chrominance
10 block in the macroblock.

The sums of differences are carried forward to the next macroblock and added to the corresponding differences (SUM[1], SUM[2],....SUM[6]). After processing each of the macroblocks of the current video frame, a summed difference is obtained for each luminance block and each chrominance block of the current video frame. Each of the six SUMS is
15 compared to its own upper and lower threshold specific to the type of block for which the SUM has been totaled. This method allows different threshold comparisons for each type of block.

If SUM[i] is greater than a predetermined threshold (thresh1[i]), in this example, where:

20 $\text{thresh1}[i] = 0.3 * \text{ABS}(\sum_{k,j} \text{DCT2}_{k,j}[i]),$

the current video frame is saved in a frame memory for further processing and possible use in the visual index. The frame memory may be a tape, a disk, as in the present invention, or any other storage medium, external or internal to the present system.

If SUM[i] is less than a predetermined threshold (thresh2[i]), where:

25 $\text{thresh2}[i] = 0.02 * \text{ABS}(\sum_{k,j} \text{DCT2}_{k,j}[i]),$ a static scene counter (SSctr) is increased to

indicate a possible static scene. The previous video frame is saved in a temporary memory.

The temporary memory only saves one frame, thus, the previous video frame will replace any video frame currently stored in temporary memory. When the counter reaches a predetermined number, (in this example, 30) the most previous video frame saved in the

30 temporary memory is transferred to the frame memory for possible use in the visual index.

Although, it is described how the next to last frame is saved to possibly represent a static scene, clearly one skilled in the art could save and use a first frame of a possible static scene in this method and the following methods.

If SUM[i] is between the two thresholds, SSctr is reset to zero and the next consecutive frames are compared.

In accordance with the instant invention the threshold(s) is (are) varied based on the "cut rate" or the amount of action in the video. This is explained in the section below
5 entitled: Determining The Threshold.

Method Two:

$$\text{SUM} = \sum_{k,j} \sum_i (\text{DCT1}_{k,j,i} - \text{DCT2}_{k,j,i})^2 / \text{ABS}(\text{DCT2}_{k,j,i})$$

- 10 k is the number of macroblocks in width of a frame, $k = 1$ to Frame-width/16,
j is the number of macroblocks in height of a frame, $j = 1$ to Frame-height/16,
i is the number of blocks in a macroblock, $i = 1$ to number of blocks in macroblock,
DCT1_{k,j,i} and DCT2_{k,j,i} are DCT coefficients for the specified macroblock and block for a
previous and a current video frame, respectively, and
15 ABS is an absolute value function.

Method two, in contrast to method one, does not discriminate between block types. Instead, method two keeps a running total of DC differences between macroblocks of current and previous video frames.

- Each difference between blocks is squared and then normalized to the DCT
20 value of the current block. Specifically, the DCT value of a block from the current video frame is subtracted from the corresponding DCT of the corresponding block in the previous video frame. The difference is then squared and divided by the corresponding DCT value of the current video frame. If the current video frame DCT value is zero, the sum for that comparison is set to one. The differences for each of the DCT values of each block in each of
25 the macroblocks of the frames are summed together to achieve a total sum, SUM.

The SUM is next compared to predetermined thresholds. If SUM is, in this example, greater than a predetermined threshold (thresh1), where:

$\text{thresh1} = 0.3 * \text{ABS}(\sum_{k,j,i} \text{DCT2}_{k,j,i})$, the current video frame is saved in the frame memory for further processing and possible use in the visual index.

- 30 If SUM is less than, in this example, a predetermined threshold (thresh2), where:
 $\text{thresh2} = 0.02 * \text{ABS}(\sum_{k,j,i} \text{DCT2}_{k,j,i})$, a static scene counter (SSctr) is increased to indicate a possible static scene. As in method one, the previous video frame is saved in a temporary memory which only saves the most previous frame. When SSctr counter reaches a

predetermined number (in this example, 30), the most previous video frame saved in the temporary memory is transferred to the frame memory for possible use in the visual index. If SUM is between the two thresholds, SSctr is reset to zero and the next consecutive frames are compared.

- 5 In accordance with the instant invention the threshold(s) is (are) varied based on the "cut rate" or the amount of action in the video. This is explained in the section below entitled: Determining The Threshold.

- 10 **Method Three:** $SUM [i] = \sum_{k,j} (DCT1_{k,j} [i] - DCT2_{k,j} [i])^2 / ABS(DCT2_{k,j})$
where:

k is the number of macroblocks in width of a frame, $k = 1$ to Frame-width/16,

j is the number of macroblocks in height of a frame, $j = 1$ to Frame-height/16,

- 15 i is the number of blocks in a macroblock, $i = 1$ to number of blocks in macroblock,
DCT1_{k,j} and DCT2_{k,j} are DCT coefficients for the specified macroblock for a previous and a current video frame, respectively, and

ABS is an absolute value function.

- Method three like method one, differentiates between each of the blocks (four
20 luminance and two chrominance blocks) of the macroblocks. In this method, DC values for each luminance and chrominance block in a current macroblock from the current video frame are respectively subtracted from the corresponding DC value for the corresponding block in the previous video frame. As in method two, however, each difference between blocks is squared and then normalized to the DCT value of the current block. Specifically, the DCT
25 value of a block from the current video frame is subtracted from the corresponding DCT of the corresponding block in the previous video frame. The difference is then squared and divided by the corresponding DCT value of the current video frame. If the current video frame DCT value is zero, the sum for that comparison is set to one.

- The differences for each of the DCT values of each type of block in each of
30 the macroblocks are summed together to achieve a total sum for the type of block, SUM[i]. Separate sums of differences, SUM [i] are kept for each of the luminance and chrominance blocks in the macroblock. The sums of differences are carried forward to the next macroblock and added to the corresponding differences (SUM[1], SUM[2],.....SUM[6]). After processing each of the macroblocks of the current video frame, a summed difference is

obtained for each luminance block and each chrominance block of the current video frame. Each of the six SUMS is compared to its own upper and lower threshold specific to the type of block for which the SUM has been totaled. This method allows different threshold comparisons for each type of block.

5 If SUM[i] is greater than a predetermined threshold (thresh1[i] as previously defined), the current video frame is saved in the frame memory for further processing and possible use in the visual index.

If SUM[i] is less than a predetermined threshold (thresh2[i] as previously defined), a static scene counter (SSctr) is increased to indicate a possible static scene. The previous video frame is saved in a temporary memory which, in the present invention, saves only the most previous video frame. When SSctr reaches a predetermined number, 30, the most previous video frame saved in the temporary memory is transferred to the frame memory for possible use in the visual index.

10 If SUM[i] is between the two thresholds, the SSctr is reset to zero and the next consecutive frames are compared.

In accordance with the instant invention the threshold(s) is (are) varied based on the "cut rate" or the amount of action in the video. This is explained in the section below entitled: Determining The Threshold.

20

Method Four:

Methods one through three each work over the complete video frame, summing either the difference or square of the difference for the DCT values for all luminance and chrominance added together or summed as individual components. Method four works on the macroblock level providing an efficient result with limited processing.

$$SUM = \sum_{k,j} Mbdiff(MB1[i]_{kj} - MB2[i]_{kj})$$

where:

k is the number of macroblocks in width of a frame, $k = 1$ to Frame-width/16,
j is the number of macroblocks in height of a frame, $j = 1$ to Frame-height/16,
30 I is the number of blocks in a macroblock, $I = 1$ to number of blocks in macroblock,

MB1_{kj} and MB2_{kj} macroblocks for a previous and a current video frame, respectively, and

Mbdiff is a function that determines the number of blocks which are different from each other between two macroblocks, and outputs a first value if this difference is higher than a certain threshold and a second value otherwise.

Specifically, a subsum (subsum[1], subsum[2],...subsum[6]) is determined for
 5 each of the blocks (Y1, Y2, Y3, Y4, Cr and Cb) of a specific macroblock by comparing a respective block of a first macroblock to a corresponding respective block of a second macroblock to obtain a subsum[i] where:

$$\text{subsum}[i]_{j,k} = \text{ABS}(\text{DCT1}[i]_{j,k} - \text{DCT2}[i]_{j,k})$$

For example, the DC value of Cr of the first macroblock of the current frame
 10 is subtracted from the DC value of Cr of the first macroblock of the previous frame to obtain a subsum[Cr]_{1,1}. Each subsum[i] is compared to a predetermined threshold (th1). If the subsum[i] is, in this example, greater than a first predetermined threshold (th1), in this example, where:

th1 = 0.3 * subsum[i], a block counter (B1ctr) is incremented and if, lower
 15 than a second predetermined threshold (th2), where:

th2 = 0.02 * subsum[i], a block counter (B2ctr) is incremented. Each
 respective subsum[i] is compared to the thresholds (th1 and th2) which may be a constant(s), based on a fixed function(s) or based on a function(s) or constant(s) specific to the type of block.

20 After the six blocks of the macroblock have been processed, the block counters are analyzed. If the block counter B1ctr is, in this example, above a predetermined threshold (B1th), in this example, three, the macroblock is considered different from the corresponding macroblock of the previous video frame and a macroblock counter, MB1ctr, is incremented. The B1ctr is then reset and a next macroblock is analyzed.

25 When all the macroblocks of a video frame have been processed, MB1ctr is compared to predetermined frame thresholds. If MB1ctr is, in this example using a 352 x 240 frame (or image), above a first predetermined frame threshold (f1th) of 100, the current frame is saved in the frame memory and MB1ctr is reset.

If some number of blocks in a macroblock are similar, B2ctr is above a
 30 predetermined threshold (B2th) of three, the macroblocks are considered the same and a second macroblock counter, MB2ctr, is incremented. B2ctr is then reset and a next macroblock is analyzed. After all the macroblocks of a frame have been analyzed, if the second macroblock counter is above a second predetermined frame threshold (f2th) of 250, the video frames are considered the same and a frame counter (Fctr) is set. MB2ctr is reset

and a next frame is analyzed. When Fctr reaches a predetermined threshold (SSthresh, in this example, 30), a static sequence is assumed to have occurred. The previous video frame is then saved in frame memory to represent the static sequence. This process continues until a video frame is determined to be different from the previous video frame or new frames are exhausted. Fctr is then reset and the next video frame is analyzed.

Those frames saved in frame memory in this and the preceding methods are considered keyframes.

Method four could also be implemented by using the normalized square of the differences. Specifically, instead of just using the difference between blocks, the difference would be squared and divided by the values found in the subtracted block. Scene cut detection may then be less sensitive to thresholds.

Keyframe filtering, discussed below, may be performed as each frame is processed under the significant scene detection process or after all the frames have been processed. Additionally, the thresholds set forth above may easily be altered to provide lesser or greater detection. For example, in 0.3 could easily be altered to any other value as could 0.02, or constants may be altered to allow for more or less efficiency, for example, SSctr could be different. Moreover, each threshold may instead be a constant, a fixed function, or a function variable on the type or location of the block being analyzed.

In accordance with the instant invention the threshold(s) is (are) varied based on the "cut rate" or the amount of action in the video. This is explained in the section below entitled: Determining The Threshold.

Determining the Threshold

As stated above, detecting where a scene cut in the video occurs is based on how much change occurs between two frames. The amount of change is measured against at least one threshold and if the change is above this threshold a determination is made that a scene change or "cut" has occurred. If the amount of change is below this threshold, a determination is made that a scene change or "cut" has not occurred. For an action movie, many times there are large differences between two frames, but a scene change has not occurred, such as a car chase scene or when an object flies through the air. Similarly in a static type of program for instance, a news broadcast which has a small box in the top right corner of the frame showing the news story, if the top right hand corner of the frame changes from one news story to another, this small change is a scene change. Therefore there is a need

for the threshold to be varied in dependence on the category of the program or movie. Figure 8a shows a video archival system which uses dynamic thresholds to determine scene changes.

Figure 9 shows a flow chart of a threshold setter which sets the threshold based on the category of the movie and compares the difference found between two frames to this threshold.

In Figure 9 the video is received (901) and the video stream is analyzed to determine whether a threshold has been sent along with the video stream (902). If the threshold is sent, the difference between two frames (903) is compared to the threshold to determine if a scene cut has occurred. If there is no threshold in the video stream then the video stream is analyzed to determine whether a category of the video (action, drama etc.) is being sent with the video stream (904). If a category is being sent, then a threshold is used that corresponds to this category (905). If a category is not being sent then the category and/or threshold is set manually at the receiver (906) based on the user's knowledge of the type of video being received. A video index is then made using the manually set threshold to detect cuts or key frames. The term "category" defines, for instance, the movie type, e.g. drama, sports, action etc., or the type of any subportions of the video e.g. high action scenes or content or static scenes or content etc.

If an index is to be made of the action movie, the threshold(s) to detect scene cuts should be set higher than if an index is to be made of the static type of broadcast.

Alternatively for a more static movie, e.g. a drama, the threshold to detect static scenes should be set lower than what would be used in an action movie. Accordingly the invention pertains to adjusting the threshold based on the category of the video. For a high action video the threshold for a scene change is set high, whereas for a news broadcast the threshold for a scene change is set lower. This varying of the threshold avoids creating a large video index for the action movie and a small video index for the news broadcast.

The thresholds for a particular video or a portion of a video can be placed in the header of the received video. Alternatively the thresholds can be sent by a third party source such as with the electronic programming guides, either for each program or for subportions of a program. By varying the threshold within a movie, it takes into account a movie which is generally a low action drama but contains high action scenes. It also avoids creating a video index which contains many high action frames and few low action frames even though there may be scene changes occurring between the low action frames and no scene changes occurring between the high action frames. A third party service could also provide the thresholds in a lookup table and transmit the thresholds to the receiver with the

incoming video, provide the thresholds in the receiver itself or download the thresholds to the receiver. In any of these cases the received video (901) can be compared to the thresholds as the thresholds are received dynamically (902, 1003) or as they are stored in memory (902, 1003.)

5 There are also many other ways the thresholds can be received, set or used which would be obvious to one skilled in the art. For example, for home movies, or movies that the user knows the content, the user can set the threshold himself. This would enable the user to create a personalized video index of home movies.

 It should also be noted that the threshold could be set very high for
10 commercials, so that many keyframes aren't selected from the commercials.

 The category of the video or portions of the video could be set by the movie maker and provided within the video stream. The receiver could also store various thresholds and select a threshold based on the category of the received video. (See Fig. 9.)

 Alternatively the video indexing system could set the threshold based on the
15 perceived content of the video determined from the movie itself (Fig. 10). To determine the threshold(s) to be used for a movie or video, or portions thereof, a universal threshold is first chosen to select key frames from the video. If the number of detected scene changes (cuts) using the universal threshold is above a certain amount (1003,1004,1005), the threshold is set higher (1006) until a reasonable number of scene changes are detected. (See Fig. 10.)

20 In Fig. 10, the video is received (1001) and two frames are compared to detect the number of differences between the two frames (1002). The number of differences is then compared to the universal threshold to see if a "cut" has occurred between the two frames (1003). The detected cuts are then counted (1004). If the number of detected cuts is too high for creation of a reasonably sized video index (1005) the universal threshold is adjusted
25 (1006) so that fewer cuts will be perceived as occurring. Once the number of detected cuts is within a reasonable range (which depends on the size of the video index that is being created) then this threshold is used (1007) to create a video index.

30 **Keyframe Filtering**

 A keyframe filtering method can also be used to reduce the number of keyframes saved in frame memory by filtering out repetitive frames and other selected types of frames. Keyframe filtering is performed by a keyframe filterer 240 in the host processor 210 after significant scene detection (Figures 2A and 2B). During the significant scene

detection process, a large number of keyframes (frames or images selected) may be extracted, for example, 2000 keyframes per hour of recorded video which is often too many to be easily handled by a user. However, from a user's perspective, not all the keyframes selected in the significant scene detection are important or necessary to convey visual contents of the video.

- 5 For example, in a dialogue scene, speakers are likely shown several times. The present invention allows retention of only one frame per speaker for the dialogue scene.

Figures 6A and 6B are an overview of the procedure for keyframe filtering. As shown in steps 602 to 606 of Figure 6A, a block signature is derived for each block in a frame. The block signature 700 is, in this example, eight bits, three of which represent a DC signature 702 and five of which represent an AC signature 704, as shown in Figure 7. All other DCT coefficients in a block besides the DC value are AC values.

The DC signature is derived by extracting the DC value (step 602) and determining where the DC value falls within a specified range of values (step 604), in this example, between -2400 and 2400. The range is divided into a preselected number of intervals as shown in Figure 7. In the present invention, eight intervals are used, although more or less intervals may be used for greater or lesser granularity of an image which can also be set by the category of the movie.

Each interval is assigned a predefined mapping such as that shown in Figure 7. Each DC value is compared to the range and the mapping for the interval into which the DC value falls is returned. The value represented by the bits needed corresponds to the number of intervals. In this example, since the range is divided into eight intervals, three bits are used. As shown in Figure 7, the block signature 700 will thus include the DC signature 702 as the first three bits and the AC signature 704 as the remaining five bits. The number of bits allocated to the DC signature or the AC signature can also be set based on the category of the movie.

In step 604 of Figure 6A, to give good representation of a range of AC values for the block, the five AC values closest to the DC values (A1 - A5) are extracted, as shown in Figure 7. In step 606, each of the five AC values is compared to a threshold (ACthresh), in this example, 200 and if the AC value is > ACthresh, a corresponding bit in the AC signature 706 is set to a predetermined value such as one, and if < or = to ACthresh, the corresponding bit is set to zero. The ACthresh can also be based on the category of the video.

The block signature 700 is thus obtained and using the block signatures, specific images or frames may be filtered out from the visual index, such as frames which are unicolor.

A quick method to filter out unicolor frames occurs between steps 602 and 604, relying only on the DC signature. Each DC signature 704 is compared and a count is kept of each specific DC signature which occurs, (step 660) i.e., each DC signature represents into which interval the DC value falls, so in this example, eight different DC signatures exist). If, for example, 90% of the blocks or 1782 blocks ($0.9 * 330 \text{ macroblocks} * 6 \text{ blocks}$) fall in the same interval (have the same DC signature), the image is considered unicolor (step 662) and the frame is discarded or filtered out from frame memory (step 664). Alternatively, separate counts could be kept for each type of block (Cr, Cb...) and each separate count compared to an associated threshold which can be based on the category of the movie.

After deriving the block signatures (Blk_sig) for each block of each macroblock in the frame, regions are determined. Regions are, in this example, two or more blocks, each block of which neighbors at least one other block in the region and which shares a similar block signature to the other blocks in the region. More blocks could be required to define a region if timing is desired to be decreased. Although each block signature of the frame could be compared to each other block signature and then counted to determine regions, the present invention may use a known technique such as a region growing process to determine regions within a frame (step 608).

During the region growing process, a region counter is kept for each region to count the number of blocks in a region (size), and is represented by 16 bits. Once the entire frame has been analyzed to find the regions, another known method may be used to find a centroid or center of each region, which, in this example, is based on an x-y axis reference (step 610). Both the x and y coordinates are extracted as CX and CY, respectively, and are represented by 16 bits. Each region is then assigned a region signature, Region(Blk_sigr, sizer, CXr, CYr) where r is a region number. The block signature for the region is determined based on a most dominant block signature as determined by the region growing process.

Based on specific criteria, in the present example, increasing size, the regions are sorted and region filtering may be performed (step 612). In this example, all but the largest three regions are filtered out. The remaining three regions are incorporated into a frame signature representing the frame. The frame signature, in the present example, is represented by 168 bits and of the form (Region1, Region2, Region3) or more specifically, (Blk_sig1, size1, CX1, CY1, Blk_sig2, size2, CX2, CY2, Blk_sig3, size3, CX3, CY3). The number of bits in the frame signature can be set by the category of the video which would also change the number of regions based on the category of the video.

As shown in Figure 6B, a frame comparison procedure compares a current (F2) and a previous (F1) frame based on their respective frame signatures (step 616). In the present example, the respective Region1s are compared, then the respective Region2s and lastly, the respective Region3s. Specifically, the block signatures of respective regions are compared according to the following:

$$FDiff = ABS(size1_{F1} - size1_{F2}) + ABS(size2_{F1} - size2_{F2}) + ABS(size3_{F1} - size3_{F2})$$

If FDiff is < 10, the frames are considered similar and further object processing is performed (step 620). If FDiff is > 10, the frames are considered different and neither frame is filtered under this procedure. Clearly what Fdiff is compared to can also be set by the category of the video.

Regions generally represent an "object" which may be an object, person, thing, etc. Object processing determines if an object shown in a region in a previous frame is in a same location or in very close proximity. In the present example, the object is a primary focus of the frame or one of the more primary focuses. For example, a video may focus on a child playing with a ball, where the ball moves about within the frame. If a user wants to limit the number of frames in the visual index such that she does not care where the object is within an image (step 622), then at this juncture, F2, the current frame, is filtered out of frame memory (step 624).

If a user cares where an object is within a frame and wishes to filter only frames having an object shown in a same or very close proximity, several methods may be used (object filter, step 626).

A first method compares centers by determining their Euclidean distances, as shown below. Specifically,

$$Edist = \sqrt{(CX1_{F1} - CX1_{F2})^2 + (CY1_{F1} - CY1_{F2})^2} + \sqrt{(CX2_{F1} - CX2_{F2})^2 + (CY2_{F1} - CY2_{F2})^2} + \sqrt{(CX3_{F1} - CX3_{F2})^2 + (CY3_{F1} - CY3_{F2})^2}$$

If Edist is > 3, the object is assumed to have moved and no filtering is performed. If Edist is < or = to 3, the object is assumed to have remained in approximately the same position and thus, the current frame is filtered out. Edist can also be compared to a number based on the category of the video.

A second method for object filtering compares frames using macroblocks. Specifically, block signatures of respective blocks within respective macroblocks are compared. For example, the block signature of the Y1 block of MB1, 1 (macroblock in

position 1,1 of a frame) of a current frame (F2) is compared to the block signature of the Y1 block of MB1,1 of a previous frame (F1).

First, the DC signature of the Y1 blocks are compared. If the DC signatures match, the AC signatures are compared, bit by bit. A count (ACcount) is kept and if a
5 preselected number of bits match, in the present example, four of five bits, a block counter (BlkCTR) is incremented. If the DC signatures do not match, or if the ACcount is < 4 , then the next block is analyzed. ACcount can be compared to a number based on the category of the video.

Each block of the macroblock (in this example using 4:2:0, six blocks) is
10 analyzed. When all the blocks in a macroblock are analyzed, the block counter is checked. If BlkCTR is 4, then the blocks are deemed similar and a macroblock counter (MBCTR) is increased. BlkCTR can also be compared to a number based on the category of the video.

Once all the macroblocks in an image have been analyzed, MBCTR is checked. If MBCTR is, in this example, $> \text{ or } =$ to 75% (247 or $0.75 * 330$ macroblocks) of
15 the macroblocks in a frame, the frames are deemed similar and the current frame (F2) is filtered out from the frame memory. If MBCTR is $< 75\%$, then no frames are filtered at this point. Again the MBCTR can be compared to a percentage which is based on the category of the video.

An additional method for filtering out unicolor frames occurs when the region
20 sizes are determined. If a region size is $\geq 90\%$ of the frame blocks or 1782 blocks, the frame is deemed to be unicolor and is filtered from frame memory. This filtering requires more processing than the previous unicolor frame filtering method discussed. Again the 90% can be varied based on the category of the video.

Based on the keyframe signature, keyframes are filtered out to retain only
25 those most likely to be desired by a user. By using different thresholds, the number of keyframes filtered out may be increased or decreased.

In the keyframe filtering process, the presence of commercials in the source video can generally be determined. The present invention allows the user to choose whether to include keyframes from the commercials as part of the visual index or instead, exclude
30 those keyframes.

Presence of commercials is generally indicated by a high number of cuts per time unit. However, action movies may also have prolonged scenes with a large number of keyframes per time unit. To have more reliable isolation of commercials, a total distribution

of the keyframes in the source video is analyzed to attempt to deduce a frequency and a likelihood of segments with commercials.

Commercials are typically spread over fixed intervals during television programs, for example, every five to ten minutes during a sitcom of 30 minutes. Duration of commercials is typically 1-2 minutes. Commercials are isolated by determining when a high number of keyframes per minute occurs. Specifically, relative times of each keyframe are compared to other keyframes. If a relatively high number of keyframes per minute occur, the threshold could be set to a high number for this area to avoid taking too many keyframes from the commercials or high action sequences.

During the significant scene detection process, when a frame is saved in frame memory as a keyframe, an associated frame number is converted into a time code or time stamp, indicating, for example, its relative time of occurrence. After every keyframe is extracted, a keyframe density is computed for the last one minute where:

$L1 = \text{Last minute keyframe density} = \text{number of keyframes in the last minute} / 1800$, and a keyframe density is computed for the last five minutes where:

$L5 = \text{Last five minute keyframe density} = \text{number of keyframes in the last five minutes} / 9000$.

If $L1 > (L2 * \text{constant})$, where constant is 3 in this example, then a potential commercial break is indicated. If a time stamp of the last keyframe of the last indicated commercial break is > 5 minutes, then a current commercial break is indicated and all the keyframes in the last one minute are filtered from frame memory. The filtering technique could be achieved by simply setting the threshold to a high value during this time period to avoid taking keyframes of commercials.

Video Retrieval: Once a video tape or file has a visual index, a user may wish to access the visual index. A video retrieval process displays the visual index to the user in a useable form. The user can browse and navigate through the visual index and fast forward to a selected point on the source tape or the MPEG file. Figure 8 details the retrieval process.

In step 802, the source video is rewound by, for example, a VCR or playback device, if required, to the location of the visual index, in this example, at the beginning of the tape. If the source video is on an MPEG file or disk, a pointer would point to the beginning of the storage location and would not need to be rewound. Similarly, other storage means would be properly set to the beginning of the visual index.

In step 804, the visual index is read by the VCR head, the computer, or other hardware device from the source video and saved into an index memory which can be any type of storage device. In step 806, a processor in, for example, the VCR retrieves keyframes from an index memory of the source video. In step 806, the retrieved keyframes are
5 processed to reduce size to, in this example, 120x80 pixels, although other frame sizes may easily be chosen automatically or manually by a user.

The processed frames are next transferred to the host processor 210 in step 808 which writes the processed keyframes to display memory and displays them in a user interface such as a computer display, television screen, etc.

10 In step 810, the source video is stopped once the video index has been read. A video indexing system or software allows keyframes to be displayed on a display, such as a computer monitor or television screen in step 812. In step 814, the visual index may be printed if the user desires. A user may also select a particular keyframe from the visual index as shown in step 816. If a user wishes to view the source video at that particular keyframe,
15 the source tape could then be automatically forwarded to a corresponding point on the source tape from where the keyframe was extracted and the source tape could thus be played (step 818). Alternatively, a counter could be displayed allowing a user to either fast forward through the source video or play the source video from the visual index to the selected key frame.

20 The present invention may also eliminate the significant scene detection processing and perform only keyframe filtering; however, processing would be significantly slower using currently and widely available processors.

An additional feature would allow a user to stop the playing of a video tape at any point and access the visual index for that video tape. This would require a memory or
25 buffer for storing the visual index when a video tape is first used during a session.

The present invention is shown using DCT coefficients; however, one may instead use representative values such as wavelength coefficients or a function which operates on a sub-area of the image to give representative values for that sub-area. This may be used in significant scene detection as well as keyframe filtering.

30 As recited in the appended claims there is provided a video system, comprising:
a receiver which receives frames of video;
a threshold memory which receives a threshold for a particular category of the video; and

a comparator which compares a first frame to a subsequent to determine the difference between the first and subsequent frame, and compares the difference to the threshold, if the difference is greater than the threshold then a scene change is deemed to have occurred between the two frames, if the difference is lower than the threshold then a scene change is deemed not to have occurred.

Optionally, if the difference is lower than the threshold for a plurality of frames the scene can be deemed a static scene. Irrespective of whether the scene is deemed a static scene, the threshold may be received from the incoming video or the threshold may be received from a third party source. In those cases where the present video system is so provided, there may be further provided a threshold setter comprising:

- a category detector for detecting the category of the video; and
- a threshold setter for setting the threshold based on the category detected. The threshold may be set at the receiver.

While the invention has been described in connection with preferred embodiments, it will be understood that modifications thereof within the principles outlined above will be evident to those skilled in the art and thus, the invention is not limited to the preferred embodiments but is intended to encompass such modifications.

CLAIMS:

1. A system for detecting significant scenes and selecting keyframes of source video, comprising:

a receiver (Fig.2A, 2B, 230) which receives source video as macroblocks including blocks of DCT coefficients and frames;

5 a summer (Fig. 8, 806) which calculates a sum for each type of block within a macroblock based on:

$$\text{SUM}[I] = \sum_{k,j} \text{ABS}(\text{DCT1}_{kj}[I] - \text{DCT2}_{kj}[I])$$

where:

k is the number of macroblocks in width of a frame, $k = 1$ to Frame-width/16,

10 j is the number of macroblocks in height of a frame, $j = 1$ to Frame-height/16,

I is the number of blocks in a macroblock, $I = 1$ to number of blocks in macroblock,

DCT1_{kj} and DCT2_{kj} are DCT coefficients for the specified macroblock for a previous and a current frame, respectively, and

ABS is an absolute value function;

15 a threshold adjustor (Fig.10, 1006) which adjusts a first threshold and a second threshold based on the category of the source video,

a first comparator (Fig.10, 1003) for comparing each SUM[i] to the first and second

thresholds and saving the current frame as a keyframe in a frame memory if SUM[i] is

greater than the first threshold, incrementing a static scene counter if less than the second and

20 saves the previous frame in a temporary memory, and resets the static scene counter otherwise; and

a second comparator (Fig.10, 1003) which compares the static scene counter to a

predetermined number and transferring the most previous video frame saved in the temporary memory to the frame memory as a keyframe.

25

2. A system for detecting significant scenes and selecting keyframes of source video, comprising:

a receiver (Fig.2A,2B, 230) which receives source video as macroblocks including blocks of DCT coefficients and frames;

a summer (Fig.8, 806) which calculates a sum for each macroblock based on:

$$\text{SUM} = \sum_{k,j} \sum_I (\text{DCT1}_{k,j,i} - \text{DCT2}_{k,j,i})^2 / \text{ABS}(\text{DCT2}_{k,j,i})$$

k is the number of macroblocks in width of a frame, k = 1 to Frame-width/16,

j is the number of macroblocks in height of a frame, j = 1 to Frame-height/16,

5 I is the number of blocks in a macroblock, I = 1 to number of blocks in macroblock,

DCT1_{k,j,i} and DCT2_{k,j,i} are DCT coefficients for the specified macroblock and block for a previous and a current video frame, respectively, and

ABS is an absolute value function;

a threshold adjustor (Fig.10, 1006) which adjusts a first threshold and a second threshold

10 based on the category of the video;

a first comparator (Fig. 10, 1003) which compares each SUM to the first and second

thresholds and saves the current frame as a keyframe in a frame memory if SUM is greater than the first threshold, increments a static scene counter if less than the second and saves the previous frame in a temporary memory, and resets the static scene counter otherwise; and

15 a second comparator (Fig.10, 1003) which compares the static scene counter to a predetermined number and transfers the most previous video frame saved in the temporary memory to the frame memory as a keyframe.

3. A system for detecting significant scenes and selecting keyframes of source
20 video, comprising:

a receiver (Fig.2A,2B, 230) which receives source video as macroblocks including blocks of DCT coefficients and frames;

a summer (Fig.8, 806) which calculates a sum for each type of block within a macroblock based on:

25
$$\text{SUM} [I] = \sum_{k,j} (\text{DCT1}_{k,j} [I] - \text{DCT2}_{k,j} [I])^2 / \text{ABS}(\text{DCT2}_{k,j})$$

where:

k is the number of macroblocks in width of a frame, k = 1 to Frame-width/16,

j is the number of macroblocks in height of a frame, j = 1 to Frame-height/16,

I is the number of blocks in a macroblock, I = 1 to number of blocks in macroblock,

30 DCT1_{k,j} and DCT2_{k,j} are DCT coefficients for the specified macroblock for a previous and a current frame, respectively, and

ABS is an absolute value function;

a threshold adjustor (Fig.10,1006) which adjusts a first threshold and a second threshold based on the category of the video;

a first comparator (Fig.10, 1003) which compares each SUM[i] to the first and second thresholds and saves the current frame as a keyframe in a frame memory if SUM[i] is greater than the first threshold, increments a static scene counter if less than the second and saves the previous frame in a temporary memory, and resets the static scene counter otherwise; and
 5 a second comparator (Fig.10, 1003) which compares the static scene counter to a predetermined number and transfers the most previous video frame saved in the temporary memory to the frame memory as a keyframe.

4. A system for detecting significant scenes and selecting keyframes of source
 10 video, comprising:

a receiver (Fig.2A,2B, 230) which receives source video as macroblocks including blocks of DCT coefficients and frames;

a summer (Fig.8, 806) which calculates a sum for each type of block within a macroblock based on:

15 $SUM = \sum_{k,j} Mbdiff(MB1[i]_{k,j} - MB2[i]_{k,j})$

where:

k is the number of macroblocks in width of a frame, $k = 1$ to Frame-width/16,

j is the number of macroblocks in height of a frame, $j = 1$ to Frame-height/16,

I is the number of blocks in a macroblock, $I = 1$ to number of blocks in macroblock,

20 MB1_{k,j} and MB2_{k,j} macroblocks for a previous and a current video frame, respectively, and Mbdiff is a function that determines the number of blocks which are different from each other between two macroblocks, and outputs a first value if this difference is higher than a certain threshold and a second value otherwise;

said summer comprising:

25 a subsum calculator (Fig.8, 806) which calculates a subsum for each block, where $subsum[i]_{j,k} = ABS(DCT1[i]_{j,k} - DCT2[i]_{j,k})$ and DCT1_{k,j} and DCT2_{k,j} are DCT coefficients for the specified macroblock;

a subsum comparator (Fig.10, 1003) which compares each subsum to a first predetermined subsum threshold and if greater than the first predetermined subsum threshold,
 30 increments a first block counter (Fig. 10, 1004) and compares each subsum to a second predetermined subsum threshold and if less than the second predetermined subsum threshold, increments a second block counter;

a block comparator (Fig.10, 1004) which compares the first block counter (Fig.10, 1003) and if above a first predetermined block threshold, increments a first macroblock counter and

resets the first block counter for analysis of a next macroblock, and compares the second block counter (Fig.10,1003) and if above a second predetermined block threshold, increments a second macroblock counter (Fig.10, 1003) and resets the second block counter (Fig.10, 1003) for analysis of a next macroblock;

- 5 a macroblock comparator (Fig.10, 1004) which compares the first macroblock counter and if the first macroblock counter is above a first predetermined frame threshold, saves the current frame in frame memory and resets the first macroblock counter for analysis of a next frame, and compares the second macroblock counter and if the second macroblock counter is above a second predetermined frame threshold, sets a frame counter (Fig.2, 230) and resets the
- 10 second macroblock counter;
- a frame comparator (Fig. 2, 230) which compares the frame counter to a predetermined scene threshold and if greater than the predetermined threshold, saves the previous frame in frame memory as a keyframe; and
- a threshold adjustor (Fig. 10, 1004) which adjusts at least one of i) the predetermined subsum
- 15 thresholds, ii) predetermined block thresholds, and iii) predetermined frame thresholds based on the category of the source video.

5. A system for filtering frames from a video source, comprising:

- a receiver (Fig.2A,2B, 230) which receives frames in a frame memory and macroblocks
- 20 comprising blocks of DCT coefficients in a temporary memory;
- a block signature (Fig. 7, 440) creator which creates a block signature for each block based on at least one signature threshold;
- a region grower (Fig. 6A, 608) which detects regions of each frame and a respective size;
- a centroid determiner (Fig. 6A, 610) which determines a centroid of each region;
- 25 a region signature creator (Fig.6A, 608) which assigns each region one of the block signatures of the blocks within the region, the respective size, and the centroid;
- a frame signature creator (Fig. 6A, 614) which creates a frame signature for a first frame, based on the region signatures; and a frame comparator which compares at least one feature
- of a first frame to at least one feature of a second frame to determine a frame difference and
- 30 compares the frame difference to at least one threshold and saves or discards the frame based on whether the frame difference is higher or lower than the threshold and wherein at least one of the signature threshold and the threshold are based in a category of the video received from the video source.

6. A system as recited in Claim 5, wherein:
each block comprises DCT coefficients including a DC value; and
the block signature comprises a DC signature representative of an interval within which the
respective DC value falls.

5

7. A system as recited in Claim 6, further comprising a unicolor filter (Fig.6A,
664) which detects unicolor frames which counts the number of blocks having the same DC
signature and if the number is greater than a prespecified count, filters out the frame, and
wherein the prespecified count is based on the category of the incoming video.

10

8. A system as recited in Claim 5, further comprising a region filterer which
filters out regions not meeting prespecified criteria, wherein the prespecified criteria is based
on the category of the incoming video.

15

9. A system as recited in Claim 5, wherein:
each block comprises DCT coefficients including a DC value; and
the block signature comprises an AC signature of bits such that a prespecified
number of bits represents whether a corresponding AC value is above a prespecified value,
and wherein at least one of the prespecified number and prespecified value are based on the
category of the incoming video.

20

10. A system as recited in Claim 5, further comprising a unicolor filter (Fig.6A,
664) which compares the corresponding size of each region within the frame to a prespecified
number and if the corresponding size is greater than the prespecified number, discarding the
frame from frame memory, and wherein the prespecified number is based on the category of
the incoming video.

25

11. A video system, comprising:
a receiver (Fig.9, 901) which receives frames of video;
a threshold memory (Fig. 9, 902) which receives a threshold for a particular
category of the video; and
a comparator (Fig. 10, 1002) which compares a first frame to a subsequent
frame to determine the difference between the first and subsequent frames, and compares the
difference to the threshold, if the difference is greater than the threshold then a scene change

30

is deemed to have occurred between the two frames, if the difference is lower than the threshold then a scene change is deemed not to have occurred.

12. The video system in accordance with claim 11 wherein if the difference is
5 lower than the threshold for a plurality of frames the scene is deemed a static scene.

13. A threshold setter, for use by a video system in accordance with claim 11 or
12, comprising:

10 a category detector (Fig.9, 904) for detecting the category of the video; and
a threshold setter (Fig. 9, 905) for setting the threshold based on the category
detected.

14. A video system as claimed in claim 11 or 12, wherein the threshold is
automatically changed if the number of detected scene changes is above a predetermined
15 number or the threshold is automatically changed if the number of detected scene changes is
below a predetermined number.

15. A filtering system for use in the video system of claim 11 or 12, which
receives key frames from each scene change and filters out at least one of i) redundant, ii)
20 similar and iii) unicolor key frames.

1/10




FIG. 1




FIG. 2A

3/10




FIG. 2B

4/10




FIG. 3




FIG. 4

5/10




FIG. 5

6/10

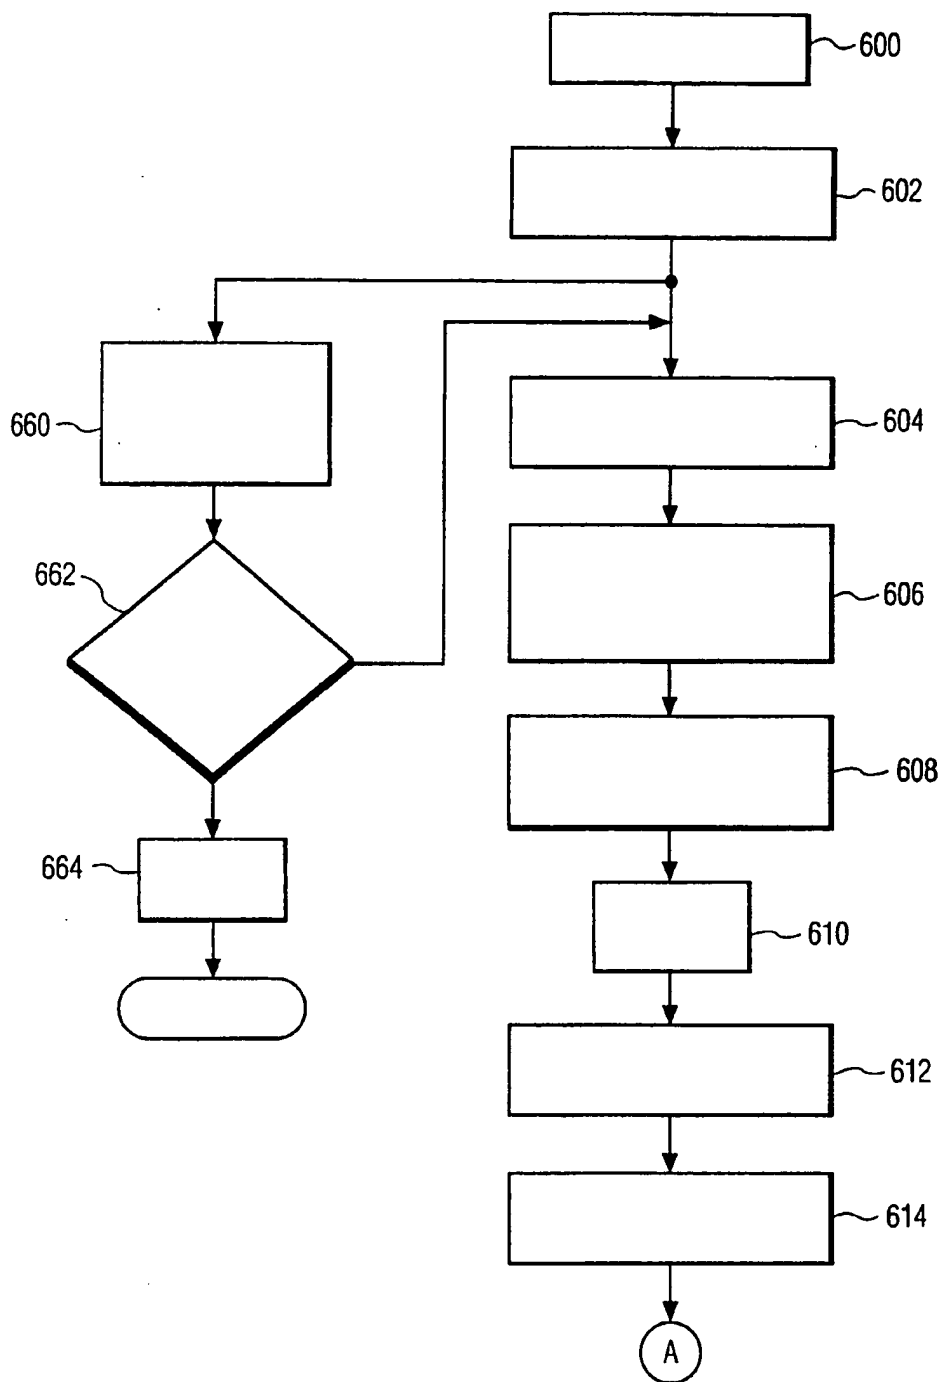


FIG. 6A

7/10




FIG. 6B

8/10




FIG. 7

9/10




FIG. 8




FIG. 8A

10/10




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/EP 00/10161

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04N7/36 H04N5/14 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 98 55943 A (KONINKL PHILIPS ELECTRONICS NV ; PHILIPS SVENSKA AB (SE)) 10 December 1998 (1998-12-10) page 1, line 26 -page 18, line 24 claims 1-3,5,6,8-11	11,12,15
Y		1-4,13, 14
A	---	5-10
	-/-	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

A document member of the same patent family

Date of the actual completion of the international search

5 February 2001

Date of mailing of the international search report

13/02/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2040, Tx: 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Fassnacht, C

INTERNATIONAL SEARCH REPORT

Int: ional Application No

PCT/EP 00/10161

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	LIENHART R ET AL: "On the detection and recognition of television commercials" PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON MULTIMEDIA COMPUTING AND SYSTEMS. LOS ALAMITOS, CA, US. 3 June 1997 (1997-06-03), pages 509-516, XP002154465 page 511, right-hand column, paragraph 3 page 515, left-hand column, paragraphs 3,4	1-4, 13, 14
A	-----	5-12, 15
A	DIMITROVA N ET AL: "VIDEO KEYFRAME EXTRACTION AND FILTERING: A KEYFRAME IS NOT A KEYFRAME TO EVERYONE" PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON INFORMATION AND KNOWLEDGE MANAGEMENT. CIKM, US, NEW YORK. ACM, vol. CONF. 6, 10 November 1997 (1997-11-10), pages 113-120, XP000775302 ISBN: 0-89791-970-X the whole document	1-15
A	----- MCGEE T ET AL: "PARSING TV PROGRAMS FOR IDENTIFICATION AND REMOVAL OF NON-STORY SEGMENTS" PROCEEDINGS OF THE SPIE, 1998, XP000972571 the whole document -----	1-15

INTERNATIONAL SEARCH REPORT

information on patent family members

International Application No

PCT/EP 00/10161

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9855943 A	10-12-1998	US 6137544 A	24-10-2000
		EP 0944874 A	29-09-1999